15-896: Algorithms, Games and Networks April 4, 2013

Lecture 19

Lecturer: Ariel Proccacia *Scribe: Arda Antikacioglu*

1 Overview

In the last lecture we didn't think of hospitals as players in a national-scale transplant game, and just focused on optimizing the number of exchanges that happened with whatever patients were registered to the exchange. In this lecture, we'll think of hospitals as selfish players who above all want to maximize the number of transplants their patients receive. This might lead hospitals to hide their patients from the system, and we'll describe the following strategyproof mechanisms to deal with issue:

- $\text{MATCH}_{\{\{1\},\{2\}\}}$ mechanism for two players.
- $\text{MATCH}_{\{\Pi_1,\Pi_2\}}$ mechanism for n-players.
- MIX-AND-MATCH mechanism for n-players.

2 Background and Motivation

In practice, hospitals care about their own track record. They want to save as many of their own patients as possible, which makes sense since those patients are lives the hospitals is directly in control of. Therefore, if a hospital can match a pair of patients without having to go through a kidney exchange, then they will go ahead and perform the transplants. There are two problems with this. The first is that this can lead to a suboptimal number of kidney transplants being performed. The second and harder to measure consequence is that easy to match transplants are repeatedly performed locally while hard to match patients stay for a long time in the kidney exchanges which degrades their quality and decreases their usefulness.

So instead of focusing on the average case behavior of kidney exchanges, we'll look at the how they can be gamed, and how we can incentivize hospitals to tell the always report all of their transplant candidates to the exchange. Unfortunately, this is a really hard to solve problem. It's so hard to We want to maximize the total number of patients that get kidneys, and make hospitals not benefit from hiding their patients. This isn't solved at all. The result we'll describe is very partial, we can't even do 3-cycles.

The model we will be working with is called the strategic model. There is an underlying graph $G = (\bigcup V_i, E)$ where each vertex represents a patient/donor pair. Edges on the graph represent compatibility. Each hospital controls a subset V_i of the patient/donor pairs and a strategy for hospital i is some set $V_i' \subseteq V_i$ of patients that it reveals to the mechanism. An edge is called internal if it's between two patients at the same hospital and external otherwise. Note that while hospitals can hide patients, they can't hide the edges between two unhidden vertices because this information can just be looked up in the medical database. A mechanism is any procedure that looks at the subgraph induced on the revealed vertices and returns a matching (so only 2-cycles are allowed). The utility of a matching to a hospital is the number of patients it matches from that hospital. Finally, a mechanism is strategyproof if revealing all the patients is a dominant strategy for all hospitals.

To see that this problem is hard, we prove note that a deterministic stratefyproof mechanism cannot do better than a 2-approximation for this problem even with only two hospitals. Consider the graph G in figure (a). Every maximum matching of G has size 3 and leaves some vertex out. This means that either the mechanism gives the grey hospital 2 matches or the white hospital 3 matches.

If the former is true, then the G' in figure (b) can only match one pair because otherwise the white hospital can hide the dashed vertices and guarantee themselves a utility of 4, contradicting strategyproofness. Similarly, if the latter is true, then G'' in figure (c) can only match one pair because otherwise the grey hospital could hide the dashed vertices and guarantee themselves a utility of 3, contradicting strategyproofness. However, OPT is size 2 for both G' and G'' which shows that strategizing can hurt the optimum. However, Ariel doesn't know any examples where one hospital's strategizing leads to a decrease in social welfare so all of them must strategize and all at the same time.

In the homework, we're asked to prove that a randomized algorithm cannot do better than 1n 8/7-approximation, but we can actually get a 6/5 lower bound with marginally more effort.

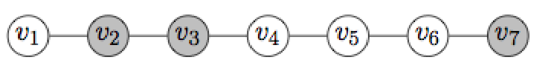

(a) Graph G. Vertices of V_1 are shown in white, vertices of V_2 in gray.

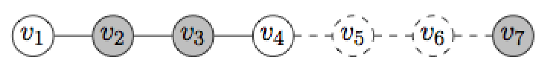

(b) Graph G' obtained from G when agent 1 hides the vertices v_5 and v_6. The hidden vertices are not part of the graph, but are shown to give the complete picture.

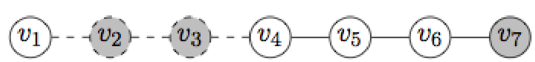

(c) Graph G'' obtained from G when agent 2 hides vertices v_2 and v_3.

3 MATCH$_{\{\{1\},\{2\}\}}$

Our first attempt at solving this problem will be a deterministic one. Consider matchings that maximize the number of "internal edges". Then among these matchings pick the one that gives the best cardinality. This mechanism is called the MATCH$_{\{\{1\},\{2\}\}}$ mechanism. We'll first prove that it's strategyproof. then that it is a 2-approximation mechanism.

Theorem 3.1. *MATCH*$_{\{\{1\},\{2\}\}}$ *is strategyproof.*

Proof. Let M be the OPT and M' the matching after player 1 manipulated his set. M' includes the hidden edges and the mechanism edges. Consider the symmetric difference $M \triangle M'$ of M and M'. Since both M and M' are matchings, their $M \triangle M'$ is a maximum degree 2 graph. These graphs are fully characterized as the disjoint union of paths and cycles. Again since M and M' are matchings, two edges that are incident on the same vertex in $M \triangle M'$ cannot both come from M

or both come from M'. This means that each path or cycle in $M \triangle M'$ is made up of alternating edges from M and M'. So cycles are of even size and have just as many edges in M and M'. This means every vertex of a cycle in $M \triangle M'$ is matched by either matching, so each player would be indifferent to the choice between the two matchings.

The only other thing $M \triangle M'$ can contain is alternating paths. So, take a path and let P be its M edges and P' be its M' edges. Denote by P_{ij} the edges of P between players i and j. Similarly, denote by P'_{ij} the edges of P between players i and j. Now we do some casework:

- Suppose first that $|P_{11}| = |P'_{11}|$. We also know that $|P_{22}| = |P'_{22}|$ since player 2 didn't hide anything. We also know that M is a maximum cardinality matching subject to being maximum cardinality on V_1 and V_2 individually, so we also have $|P_{12}| \geq |P'_{12}|$ for the remaining edges. Putting it all together shows that M is at least as good as M' for player 1:

$$u_1(P) = 2|P_{11}| + |P_{12}| \geq 2|P'_{11}| + |P'_{12}| = u_1(P')$$

- Or we can have $|P_{11}| > |P'_{11}|$ (alternatively $|P_{11}| \geq |P'_{11}| + 1$) . We claim that $|P_{12}| \geq P'_{12} - 2$. To see this, we first decompose the path into subpaths $P_1, \ldots, P_n$ which are alternately contained in V_1 and V_2 and connected by edges that cross the different hospital sets. Any such subpath completely contained in V_2 must have even length. Otherwise we can switch to the edges of one of the matchings and make get a matching that's higher cardinality on V_2. But this is impossible since player 2 is gave perfect information and the mechanism optimized for the number of matches in V_2. This means that every time we switch back and forth between V_1 and V_2, whenever we enter V_2 with an edge in M, we leave with M' and vice versa. This pairs up the vertices of M and M'. The only difference can be in the first and last edges where we may lose 2 edges, i.e. $|P_{12}| \geq P'_{12} - 2$. Once again M is at least as good as M' for player 1:

$$u_1(P) = 2|P_{11}| + |P_{12}| \geq 2(|P'_{11}| + 1) + |P'_{12}| - 2 = u_1(P')$$

Since in each case player 1 is just as well off by revealing all patients, the mechanism is strategyproof. □

Theorem 3.2. *$MATCH_{\{\{1\},\{2\}\}}$ is a 2-approximation mechanism.*

Proof. Any matching by the mechanism is inclusion maximal. Given any maximum matching on G and any edge in that matching, an inclusion maximal matching must match one of the endpoints of that edge or it wouldn't be maximal. So there are at most 2 vertices in a maximum matching for each vertex in a maximal matching and we have a 2-approximation as claimed. □

Since we proved above that a deterministic mechanism cannot give an approximation guarantee better than a factor of 2, this is the limit of what can be done deterministically. The algorithm generalizes naturally to n-players and the approximation guarantee is still valid for any number of players. However, if there are more than 2 players, then the mechanism is not necessarily strategyproof so we need a different idea to go on.

This motivates the MATCH_Π which is a different generalization of the $\text{MATCH}_{\{1\},\{2\}}$ mechanism. With this new mechanism, we fix a partition $\Pi = (\Pi_1, \Pi_2)$ be a bipartition of the players. The mechanism then rules out external edges between different players on the same side of the partition and returns a maximal cardinality matching from among the matchings that match a maximal number of patients at each hospital internally. This turns out to be strategyproof for any partitioning of the players. However, there's no approximation guarantee. This can easily be seen from the fact that all players can end up on the same side of the partition and that there might be no internal transplants to be made at any of the hospitals.

4 MIX-AND-MATCH

To solve the problems mentioned above, we introduce a mechanism called MIX-AND-MATCH. This new mechanism is strategyproof and gives a 2-approximation, but is randomized. It works by picking a random partition Π of the players and running MATCH_Π. Note we're not using randomness here to achieve strategyproofness. As we mentioned above, regardless of the partition we choose, MATCH_Π is strategy proof already. We're introducing randomness to fix the approximation guarantee of MATCH_Π. So a hospital isn't just making its expected utility better by being honest, it's making its actual utility in every case better.

Theorem 4.1. *MIX-AND-MATCH is a 2-approximation mechanism.*

Proof. Let M^* be the optimal matching and M^{**} be the union of the maximum cardinality matchings of $V_1, \ldots, V_n$ using only internal edges. We will create a new matching M' as follows. For each path P in $M^* \triangle M^{**}$, if M^{**} has more internal edges in P then we add $M^{**} \cap P$ to M'. Otherwise we add $M^{**} \cap P$ to M'. As with $\text{MATCH}_{\{\{1\},\{2\}\}}$, we don't need to worry about cycles. With this construction, M' has the following properties:

- M' is maximum cardinality on each V_i individually.
- For every internal edge M' gains relative to the optimal matching, it loses at most 2 external edges. For a matching M, we use the notation M_{ij} to denote $M \cap G[V_1, V_2]$. Using this notation:

$$\sum_i |M'_{ii}| - |M^*_{ii}| \geq \frac{1}{2} \sum_{i \neq j} |M^*_{ij}| - |M'_{ij}|$$

$$\sum_i |M'_{ii}| + \frac{1}{2} \sum_{i \neq j} |M'_{ij}| \geq \sum_i |M^*_{ii}| + \frac{1}{2} \sum_{i \neq j} |M^*_{ij}|$$

Now we fix Π and let M^π be the matching returned by $MATCH_\Pi$. Since our mechanism returns the highest cardinality matching on the given graph subject to returning a maximum cardinality matching for each hospital individually:

$$\sum_i |M^\Pi_{ii}| + \sum_{i \in \Pi_1, j \in \Pi_2} |M^\Pi_{ij}| > \sum_i |M'_{ii}| + \sum_{i \in \Pi_1, j \in \Pi_2} |M'_{ij}|$$

Finally, we can calculate the expected size of the matching M^π returned by the mechanism. Noting that all external edges must be crossing the partition:

$$
\begin{aligned}
\mathbf{E}[|M|^{\Pi}] &= \frac{1}{2^n}\sum_{\Pi}\left(\sum_i |M_{ii}^{\Pi}| + \sum_{i\in\Pi_1, j\in\Pi_2} |M_{ij}^{\Pi}|\right) \\
&\geq \frac{1}{2^n}\sum_{\Pi}\left(\sum_i |M'_{ii}| + \sum_{i\in\Pi_1, j\in\Pi_2} |M'_{ij}|\right) \\
&\geq \sum_i |M'_{ii}| + \frac{1}{2^n}\sum_{\Pi}\sum_{i\in\Pi_1, j\in\Pi_2} |M'_{ij}| \\
&\geq \sum_i |M'_{ii}| + \frac{1}{2}\sum_{i\neq j} |M'_{ij}| \\
&\geq \sum_i |M^*_{ii}| + \frac{1}{2}\sum_{i\neq j} |M^*_{ij}| \\
&\geq \frac{1}{2}\sum_i |M^*_{ii}| + \frac{1}{2}\sum_{i\neq j} |M^*_{ij}| = \frac{1}{2}|M^*|
\end{aligned}
$$

where we went from line 3 to 4 by noting that an edge crosses a random partition with probability 1/2 and all the other inequalities follow from ones mentioned above. This completes the proof. □

Unlike with the deterministic case, we didn't get a tight approximation guarantee and there's still quite a gap from 6/5 to 2 to close up.